Using a Plane-Wave Signal Model to Suppress Airborne GMTI Radar Clutter and Calibrate the Array

G.R. Legters
Science Applications International Corporation,
688 North Hedgecock Square, Satellite Beach, FL 32937 USA
legtersg@saic.com, (321)777-0061

Abstract — The prior knowledge that the radar return signal is a superposition of near-ideal plane-waves is used to suppress clutter and calibrate the array. A previous [1] radar version of the CLEAN [2] algorithm is modified to perform a parameter fit that analyzes signal data into (nonorthogonal) plane-wave modes. Signal data models now require fewer (less-fragmented) plane-wave clutter and target modes. Each range gate is processed independently and no sample covariance matrices are necessary.

INTRODUCTION

Early radar systems relied on geometries that purposely avoided the illumination of clutter or reception of interference. Target (airplane or ship) return signals were detected in relatively stationary and uncolored receiver noise. More ambitious radar systems must handle clutter and interference signals which are nonstationary and colored in some combination of range, time, space, Doppler, and angle. Whitening filters, Constant-False-Alarm-Rate (CFAR) processing, and nonstationary heuristic tricks are traditionally employed to effectively flatten the clutter and interference into stationary white noise.

Airborne Ground-Moving-Target-Indicator (GMTI) radar systems [3,4] must find targets moving on an enormous background of clutter which is also moving at a wide range of velocities relative to the receive array. This is analogous to finding needles ambling through an exploding haystack. This difficult job requires a significant space-time aperture to separate the moving targets from the clutter. The combination of a significant space-time aperture and ordinary terrain produces a highly energetic, highly nonstationary, and highly colored clutter return signal. Clutter discretes over 50 dB stronger than the targets are not unusual.

Traditional Space-Time-Adaptive-Processing (STAP) [3,4] often relies on heuristic techniques to obtain a good estimate of the covariance matrix associated with each beam-Doppler target cell. Performance depends on how accurately the covariance matrix represents the local clutter and, to some extent, on how well the steering vectors match the radar system. Larger space-time apertures theoretically allow better performance because larger covariance matrices have more degrees of freedom to cancel clutter. Unfortunately, larger apertures resolve lumpier clutter so that sample support homogeneity suffers. Reduced rank techniques ideally allow reduced sample support. Unfortunately, real-world, large-aperture, sample covariance matrices are often of significantly high rank. Because the clutter power is so much stronger than the target power, small timing or pointing perturbations of the scattering and measurement processes smear and modulate significant power from large eigenmodes into smaller eigenmodes.

It appears that traditional STAP is pushed beyond the limit when large space-time apertures observe targets buried in extremely strong clutter and interference. But airborne GMTI radar systems are trending towards larger space-time apertures to handle target-rich, highly inhomogeneous urban and mountainous environments. One can extend the use of STAP by obtaining more accurate covariance matrices based on other measurements and prior information. However, as space-time apertures increase, it becomes much more difficult to patch covariance matrix estimates.

The basic problem is that large-aperture space-time data sets have too much structure and contain too much information to permit accurate covariance matrix estimation. But highly-structured, information-rich data should present an opportunity, not an obstacle. Maybe there is a better way to extract information from the data.

Note that any multidimensional Kalman [5] filtering problem could also be solved (poorly) using STAP-like techniques. Both approaches estimate a state covariance matrix and use the equivalent of a steering vector to provide a state estimate from measurements. The Kalman filter has the advantage that known state biases (evolution matrix), measurement biases (measurement matrix), correlated process perturbations, and correlated measurement uncertainties are “removed” from the state covariance matrix. Iterative Kalman filters easily handle known nonstationary system and measurement structure.

Although the Kalman filter may not directly apply to the airborne GMTI radar problem, the overall Kalman philosophy does. In this era of higher-resolution space-time radar measurements, the known structure of the measurements should be exploited, starting with a physical model of the scattering process.

PLANE-WAVE MODEL OF GMTI RADAR DATA

Within a given range gate, an airborne GMTI radar typically observes a long, thin, curved strip of the earth's surface. Because the propagation and scattering processes obey (very nearly) linear equations, the return signal can be modeled as a linear combination of Huygens wavefronts emanating from clutter and target "point" scatterers along the observed strip. If the radar signal wavelength is λ, the "points" can be thought of as blobs of $\lambda/2$ extent, with little loss of accuracy. Because the antenna array length, D, is very much smaller than the distance, R, from the illuminated strips, scatterers are smeared into resolution cells of extent $R\cdot\lambda/(2\cdot D)$. In the neighborhood of the antenna array, the reflected signal from a resolution cell is (very nearly) in the form of a plane wave.

Over the coherent processing interval (CPI), M complex time samples are collected for each of N spatial channels. Suppose the space and time samples are, at least approximately, uniformly spaced Δx and Δt apart respectively. The measured signal at the (n,m)th space-time sample can be modeled as

$$y_{n,m}(\mathbf{A},\boldsymbol{\kappa},\boldsymbol{\omega}) \equiv \sum_{p=1}^{P} A_p \cdot \exp\left[j\cdot\left(\kappa_p \cdot n + \omega_p \cdot m\right)\right] + v_{n,m} \qquad (1)$$

$$\text{with } \kappa_p \equiv \frac{2\cdot\pi}{\lambda}\cdot\sin\left(\theta_p\right)\cdot\Delta x \quad \text{and} \quad \omega_p \equiv 2\cdot\pi\cdot f_p\cdot\Delta t \qquad (2)$$

where A_p is the complex amplitude of the pth structured (target or clutter) signal component, κ_p is the spatial frequency associated with a plane wavefront arriving at a beam angle, θ_p, relative to broadside, ω_p is the temporal frequency associated with a Doppler frequency, f_p, and $v_{n,m}$ is the measurement noise. Jammers can be included in this model. A point jammer is modeled with M or fewer strong "Doppler" frequencies from a single beam direction. Powerful ground-illuminating jammers can occupy more than one beam direction.

The redefined problem is to separately estimate the target and clutter sinusoids of (1), not to cancel clutter. For the purpose of a catchy acronym, this unconventional approach might be called Signal and Clutter as Highly Independent Structured Modes (SCHISM [1]). The method is founded on the prior information that the typical airborne GMTI radar return signal is derived from a superposition of near-ideal EM plane waves. Many of these plane waves possess a very high signal-to-random-noise ratio (S/N). Fortunately, it is possible precisely estimate the frequency and amplitude [6] of the troublesome strong clutter signals.

The coherent analysis of the radar return signal into precise clutter and signal modes provides many potential benefits. Fine beam-Doppler resolution yields accurate target positions. A precise, coherent clutter model provides a very sharp, selective interference+jammer filter through simple subtraction. An accurate estimate of the ideal clutter-only return signal can be synthesized for calibration on clutter. Non-moving scatterers are not filtered out so that the radar operates simultaneously in Moving-Target-Indicator (MTI) and Synthetic-Aperture-Radar (SAR) modes. Although a single CPI spans a rather puny synthetic aperture, multiple CPIs can be merged by use of the clutter estimate as a reference for motion-compensation and other biases. A precise estimate of the clutter modes can provide control points for registering targets to ground coordinates or for navigation. Amplitude and frequency parameters are independently estimated for each range gate so that covariance matrices are avoided (perhaps entirely). Finally, initial experiments indicate that the processing load will be lightened relative to Minimum-Variance-Distortionless-Response (MVDR) STAP or multi-beam SAR algorithms.

The catch is that fitting the model of (1) to the data requires nonlinear operations for frequency estimation and the solutions are not unique. The simple example

$$s(t) \equiv a\cdot\exp(j\cdot\omega\cdot t) \qquad (3)$$

with the solution

$$\omega = \text{angle}(s(2)/s(1)), \quad a \equiv s(1)\cdot\exp(-j\cdot\omega) \qquad (4)$$

demonstrates that the amplitude and frequency of an isolated noiseless sinusoid can exactly estimated from two (Nyquist) samples. When the frequency is known, the amplitude can be found using a linear estimator. Whether or not the amplitude is known, the frequency requires a nonlinear estimator. Naturally, it is possible to fit many functions through two data points. Fig. 1 plots a single red frequency and a linear combination of 32 blue frequencies that both precisely fit two samples of the sinusoid, $2\cdot\exp(j\cdot 4\cdot t)$. Without additional information, it is impossible to choose between the two solutions. The radar might be illuminating a single metal fence post, or a linear array of 32 fence posts. The two-sample array cannot resolve individual posts.

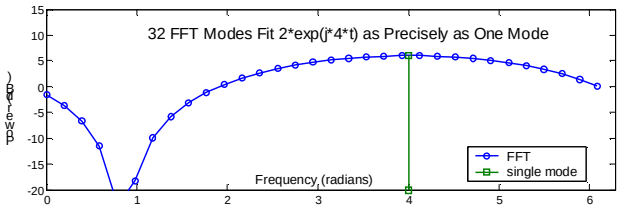

Fig. 1. Different models precisely fit two data samples.

CHOOSING A PLANE-WAVE ANALYSIS ALGORITHM

Because there is no unique solution for (1), there is also no unique algorithm to estimate $(\mathbf{A},\boldsymbol{\kappa},\boldsymbol{\omega})$ from space-time data. A (linear) fast Fourier transform (FFT) can be used to transform the space-time data into the beam-Doppler domain. The FFT precisely fits the data to an arbitrarily chosen uniform 2D grid of beam-Doppler frequencies. The result is a solution where power is spread over broad main beams and leaked into sidelobes (see the blue portion of Fig. 1). The FFT is definitely not a good candidate analysis algorithm.

A more compact set of beam-Doppler modes can be estimated using many quite different nonlinear techniques. The choice of an algorithm depends on the desired properties of the solution. The airborne GMTI plane-wave analysis algorithm should keep the number of scatterers near Brennan's rule [2,3]. It is easier to describe a diffuse scatterer as combination of point (resolution cell) scatterers than vice versa. Algorithms based on mode orthogonality do not match the model of (1) because the actual scatterer frequencies are not likely to be orthogonal over the sample support. Post-Doppler or post-beam algorithms should be avoided because true 2D algorithms exhibit better S/N and plane-wave mode isolation. However, an efficient FFT preprocessing of the space-time data to the 2D beam-Doppler mode domain can be used to localize mode power. Data tapering can reduce mode cross coupling. Performance metrics should be Minimum-Doppler-Velocity (MDV), and detection and false alarm probabilities. If the radar system reliably finds targets, it doesn't matter how well (1) fits the data.

After some experimentation, it was found that algorithms based on quadratic metrics tend to evenly distribute power among modes. For example, the blue curve of Fig. 1 can be thought of as a Least-Squared-Error (LSE) fit of the data to 32 frequency modes. Even if frequencies are poorly chosen, an LSE algorithm will try to find a way to spread out the power. A nonquadratic metric encouraged the concentration of power into compact modes, but also exhibited a capture zone limit and greatly reduced processing efficiency. An alternating projection [7] attempt required too much iteration to sharpen the modes. The Prony algorithm [8] was designed for deterministic "superresolution". However, frequency estimation algorithms based on autoregressive signal models or noise subspaces [9] do not constrain the modes to lie on the unit torus (2D version of unit circle). Also, Prony-style algorithms have issues with process order estimation and the complexity of 2D autoregressive models [10] (not Toeplitz). A previous [1] airborne GMTI radar version of the CLEAN [2] algorithm did a good job concentrating power into compact regions of beam-Doppler space. Unfortunately, the number of required modes was at least an order of magnitude higher than predicted by Brennan's rule. Algorithmic modifications to reduce the number of modes proved successful, as shown later.

GMTI Radar Version of the CLEAN Algorithm

The CLEAN algorithm [2] is often used in radio astronomy imaging to detect small radio sources obscured by nearby large sources. If the impulse response, point-spread function, Green's function, etc. is known for a measurement process, it is possible to precisely subtract out the effect of a large point interference source, which can obscure small sources. Usually, extended interfering sources can be approximated by a group of point sources.

The coherent beam-Doppler point spread function of an airborne GMTI radar system can be accurately estimated because the incident radiation is in the form of near-ideal plane waves. The CLEAN-like algorithm first finds the 2D frequency and amplitude of the largest peak of the beam-Doppler FFT of the data and then subtracts out a portion of the mode. The process is repeated until the power drops below a desired threshold.

Fig. 2 shows the operation of the CLEAN algorithm for an isolated, unresolved pair of 1D modes. The red and blue dashed vertical lines mark the amplitude and frequency of the two unresolved modes that sum to the black curve. The upper right graph shows the solid blue error residue after the first four green CLEAN modes are subtracted. The lower left graph shows the typical error residue lobes symmetrically distributed about the first mode. The lower right graphs shows that 16 green modes beat the solid blue error residue down to –50 dB. The CLEAN process represents the black sum signal as 16

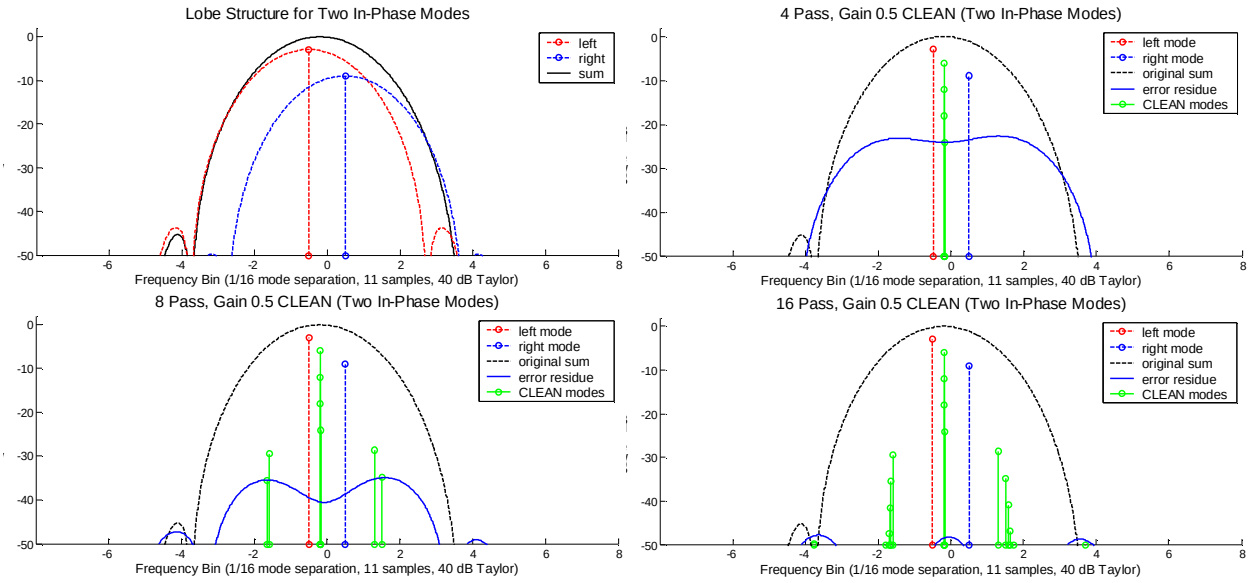

Fig. 2. Operation of the CLEAN algorithm for an isolated, unresolved pair of 1D modes.

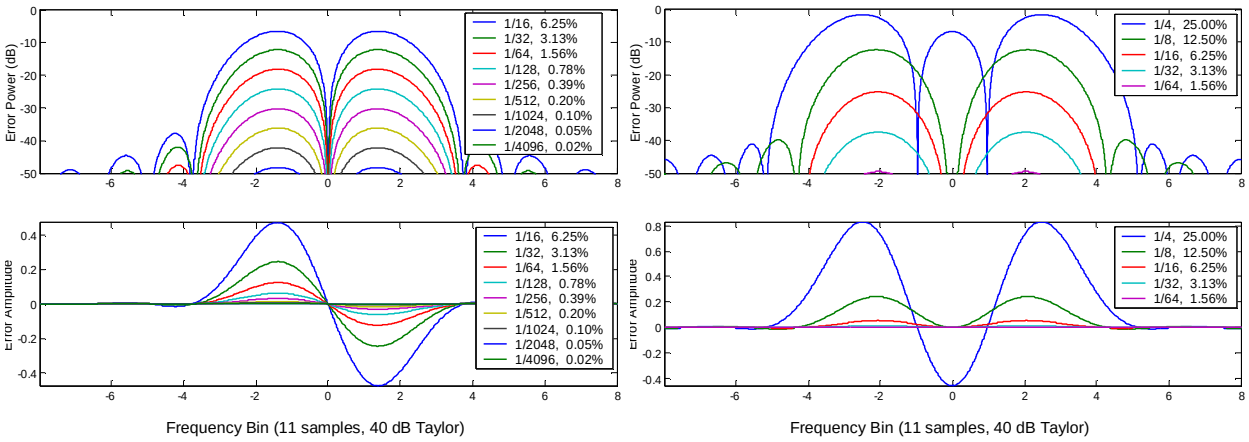

Fig. 3. Error sensitivity of mode subtraction (frequency error on left, unresolved pair error on right).

modes arranged in three clumps. A gain of 0.5 was used, which means that only half of the peak error amplitude was subtracted at each iteration. If the entire amplitude were subtracted (unity gain), the symmetric error lobes would be much more powerful. For GMTI radar, the error lobes degrade the MDV and can produce false alarms.

A heavy data taper can be used to suppress sidelobes so that amplitude estimates are not corrupted by mode cross-talk and parallel peak identification techniques do not accidentally identify a sidelobe as a mainlobe.

The left side of Fig. 3 plots the error lobes associated with frequency estimation error for an isolated mode. The top graph plots power (dB) with respect to the peak mode power, and the bottom graph shows the antisymmetric amplitude response. A fraction of 1/16 in the legend means that the estimated frequency was in error by 1/16 of the horizontal plot extent. Accuracy to 1/4096 of the bandwidth is necessary to beat the error lobes down to –50 dB. The bad news is that the model fidelity is extremely sensitive to frequency estimation error. The good news is that this sensitivity should permit very fine beamsplitting.

The right side of Fig. 3 plots the error lobes associated with modeling two unresolved modes as a single mode. The top graph plots power (dB) with respect to the peak mode power, and the bottom graph shows the symmetric amplitude response. The top curve forms a triplet because the pair separation is wide enough to resolve the two peaks (sum sags in the middle of the plot). A fraction of 1/16 in the legend means that the unresolved mode frequencies are separated by 1/16 of the horizontal plot extent. Accuracy to 1/64 of the bandwidth is necessary to beat the error lobes down to –50 dB. The insensitivity to this type of error is explained by the top left graph of Fig. 2. The sum of two unresolved peaks looks very much like a single peak.

Figure 4 shows the range-averaged beam-Doppler power for the raw KASSPER Challenge Datacube [11]. The datacube consists of 32 time samples for each of 11 spatial channels for each of 1000 range gates. The plot is generated using a 2D FFT to transform zero-filled space-time data to beam-Doppler space, calculating the beam-

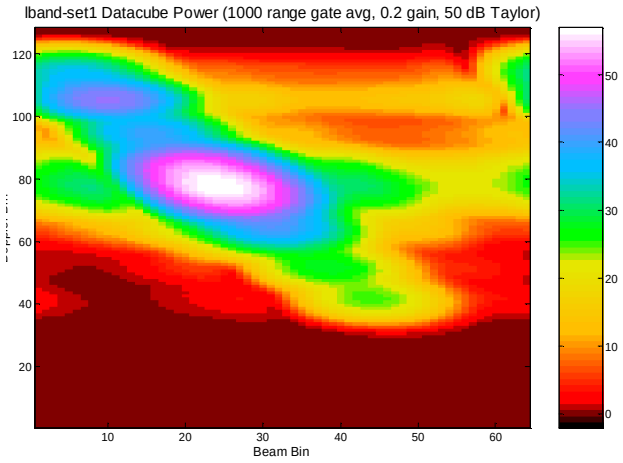

Fig. 4. Range-averaged beam-Doppler power for the raw KASSPER Challenge Datacube [11].

Doppler power for each range gate, and then averaging power over the 1000 range gates. The lowest average beam-Doppler bin is used to set the noise floor. The clutter ridge is rotated 90° from the usual convention, because the array elements happened to be arranged "backwards" in the datacube. Again, the plot displays the raw datacube, ignoring calibrated steering vectors. A 50 dB Taylor taper was used to suppress sidelobes. Note that measurement biases (array calibration errors) smear some of the clutter power off the ridge.

Fig. 5 displays the range-averaged beam-Doppler CLEAN mode power for the Challenge Datacube. A gain of 0.2 and a 50 dB Taylor taper were used to model modes greater than 13 dB above the noise floor. An unoptimized algorithm required 778 ms (Athalon XP1900) per range gate using an average of 321 modes. Each range gate consists of 352 space-time samples, and the beam-Doppler display shows that the power is concentrated in a small subset of modes. The average 321 modes per gate is an order of magnitude greater than desired. Note that measurement biases leak power into beam sidelobes near Doppler bin 300.

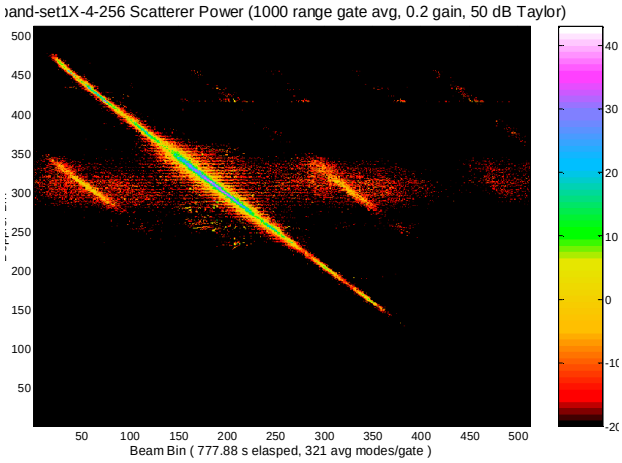

Fig. 5. Range-averaged beam-Doppler CLEAN mode power for the raw Challenge Datacube.

The scale of Fig. 4 reaches 55 dB while the scale of Fig. 5 only reaches 43 dB. The FFT modes of Fig. 4 are orthogonal, while the CLEAN modes are not. Fig. 6 clumps the modes of Fig. 5 into a 16x32 grid that is near the 11x32 degrees of freedom per range gate, restoring the maximum scale to near 55 dB. The mode plots are not a substitute for a power spectrum.

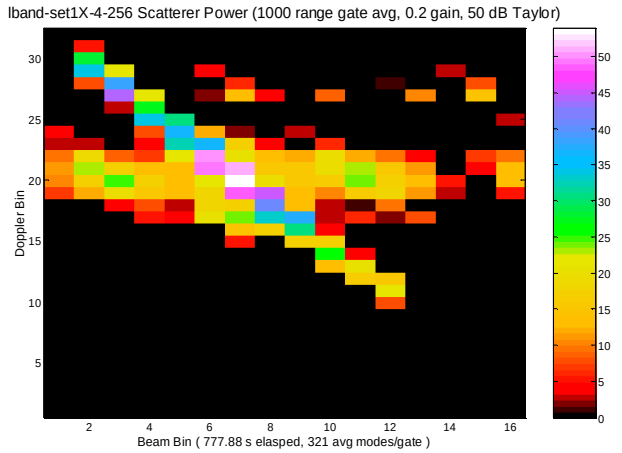

Fig. 6. Range-averaged beam-Doppler CLEAN modes coherently clumped into a 16x32 grid.

Fig. 7 displays the range-averaged estimation error beam-Doppler power for the raw KASSPER challenge datacube. The mode analysis was stopped at 13 dB above the noise floor, leaving a residual clutter ridge clearly visible. More modes could be subtracted to provide a more precise model.

The CLEAN mode analysis does a good job of finding the approximate frequencies of the beam-Doppler clutter ridge, revealing structure obscured by the FFT mainlobes and sidelobes of Fig. 4. The modes fit the data precisely to within 13 dB above the noise floor, but do not necessarily correspond with superresolved scatterer positions. The CLEAN analysis exposes the array calibration error lobes in Fig. 5.

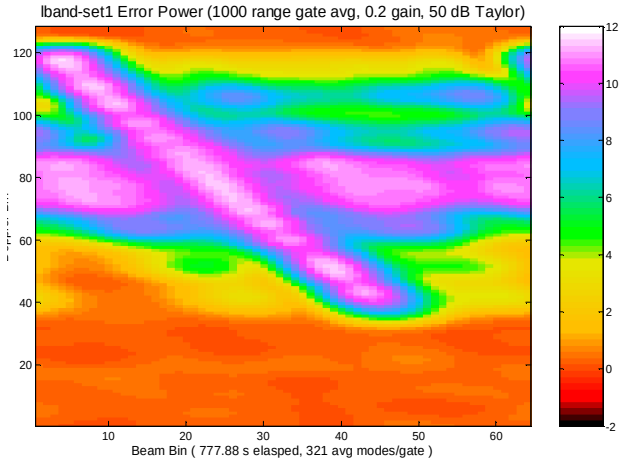

Fig. 7. Range-averaged estimation error beam-Doppler power for the raw Challenge Datacube.

CALIBRATION ON CLUTTER

Airborne GMTI radar signals are so dominated by ground clutter that expensive space-time measurements are required to reveal moving targets. Mediocre airborne radar antenna arrays cost a fortune. The cost rises dramatically when sidelobe suppression requirements exceed 30 dB. Fortunately, the clutter return provides a very strong, steady, highly-structured calibration signal. The ability to use ground clutter for constant array calibration can dramatically reduce system costs and operational constraints. The SCHISM approach was originally developed for calibration on clutter.

The modes obtained from coherent SCHISM analysis are tailor-made for calibration on clutter. The beam-Doppler ridge for the Challenge Datacube should lie along a single, nearly straight line. Fig. 5 shows error lobes on either side of the clutter ridge, which could mask or be mistaken for moving targets. A perfectly calibrated system would not exhibit the error lobes of Fig. 5. Because almost all of the average radar return power is ground clutter, an ideal sampled signal can be approximated by transforming only the clutter ridge modes of Fig. 5 back into space-time data. The array can be calibrated by comparing the ideal sampled signal to the actual sample system.

The calibration procedure can be used to find steering vectors. Each direction must separately calibrated. Alternatively, the datacube can be focused, so that the sample gain is uniform and the samples are uniformly spaced in space-time. Then, the steering vectors correspond to (weighted or unweighted) FFT modes. Datacube focusing is convenient for research purposes.

The procedure begins by fitting the dominant average clutter ridge to a quadratic. The clutter ridge is slightly curved because of the crab angle. Then, an ideal clutter model is obtained by throwing out modes not near the main clutter ridge. The rejected modes contain negligible

power. The datacube and clutter model are transformed to the range-Doppler domain. The Doppler phase of the ideal signal is compared the to the actual signal for all spatial elements. The average linear drift and offset of the datacube phase is adjusted to match the ideal signal. The datacube is transformed back to space-time, yielding a datacube with equal element spacing and calibrated channel phase. The channel amplitudes are scaled to equalize the average datacube channel power. A similar procedure can equalize time samples (not often necessary). More CPIs and adaptive processing might allow higher-order corrections, such as element pattern calibration or longer-term drift.

Fig. 8 shows the range-averaged beam-Doppler power for the focused KASSPER Challenge Datacube. The clutter ridge is less smeared compared to Fig. 4.

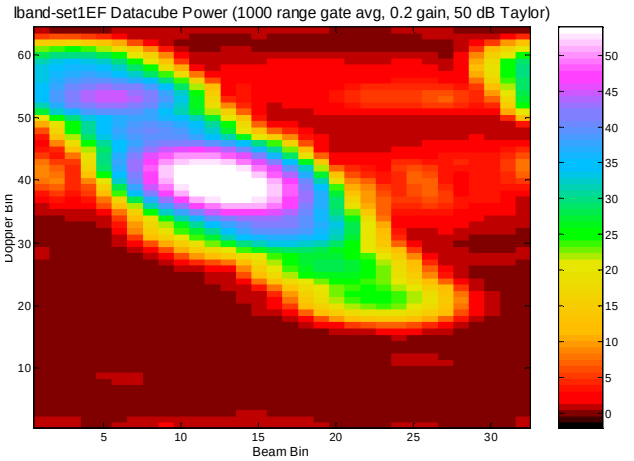

Fig. 8. Range-averaged beam-Doppler power for the focused KASSPER Challenge Datacube.

Fig. 9 shows the range-averaged beam-Doppler CLEAN mode power for the focused Challenge Datacube. The Fig. 5 measurement bias sidelobes are suppressed. The leakage from the clutter ridge is exaggerated by the 1000 range gate average. The target dots occupy one or two range gates, while clutter leakage occupies many gates.

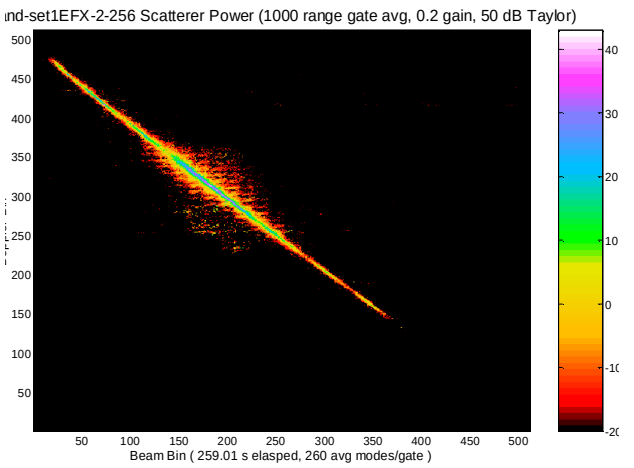

Fig. 9. Range-averaged beam-Doppler CLEAN mode power for the focused Challenge Datacube.

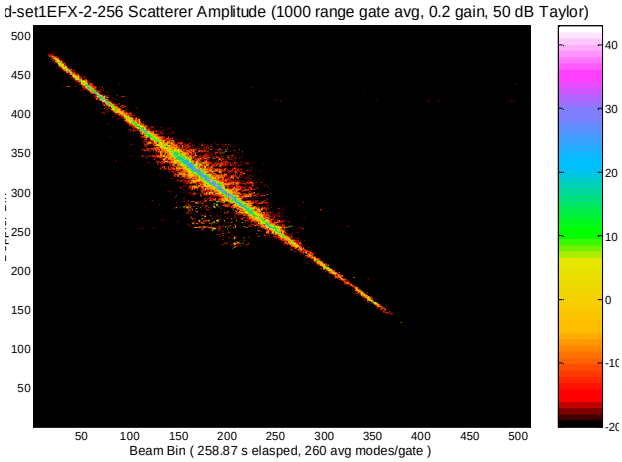

Fig. 10. Range-averaged beam-Doppler CLEAN coherent mode sum power for the focused datacube.

Fig. 10 displays the range-averaged beam-Doppler CLEAN coherent mode sum power for the focused Challenge Datacube. The modes are coherently summed to show that the error power sum exaggerates false alarm problem. The clutter ridge leakage is less noticeable.

Fig. 11 displays the range-averaged beam-Doppler CLEAN mode power for the focused datacube in more detail. Some of the bright points away from the clutter ridge correspond to targets and some to false alarms. Because the clutter is so much stronger than the targets, small model errors can produce many false alarms.

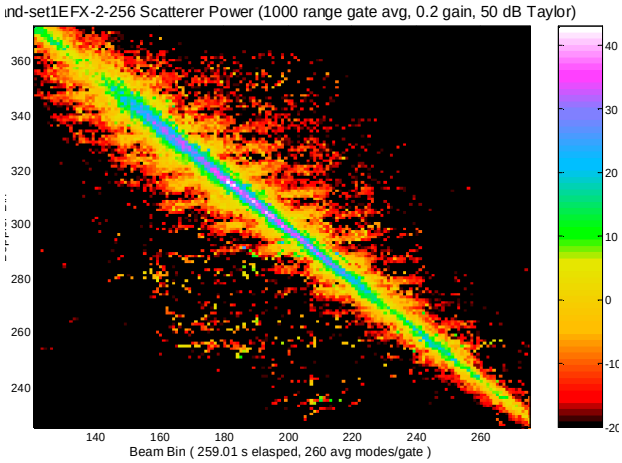

Fig. 11. Detail of range-averaged beam-Doppler CLEAN mode power for the focused datacube.

Fig. 12 shows the target-only range-averaged beam-Doppler CLEAN mode power that may be compared to Fig. 11. Again, clutter leakage is exaggerated by the 1000 range bin average. The high-fidelity simulation used to generate the KASSPER Challenge Datacube made this plot possible. The ability to decompose the datacube into target-only and targetless portions is invaluable when developing and testing algorithms. Note that the CLEAN analysis process smeared some of the targets. Fig. 2 demonstrates how the smearing can occur.

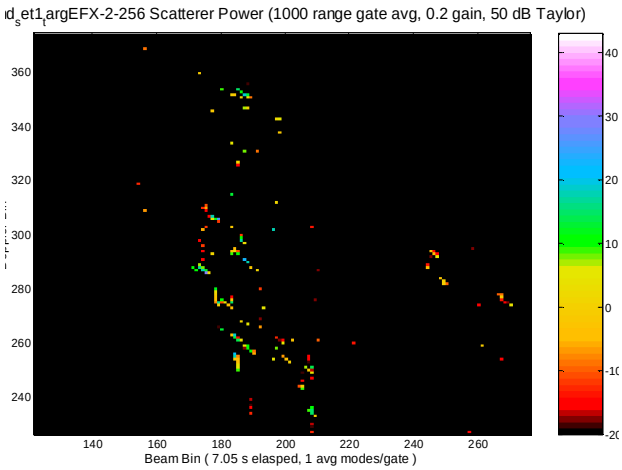

Fig. 12. Target-only range-averaged beam-Doppler CLEAN mode power.

Fig. 13 displays the range-averaged estimation error beam-Doppler power for the raw KASSPER challenge datacube. The mode analysis again stopped at 13 dB above the noise floor, leaving a residual clutter ridge clearly visible.

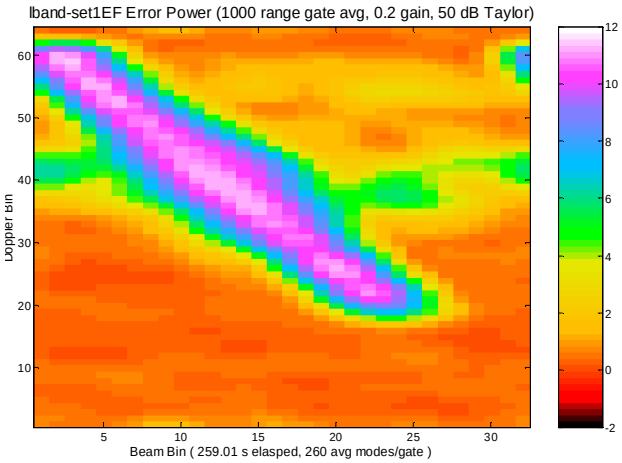

Fig. 13. Range-averaged estimation error beam-Doppler power for the focused datacube.

For completeness, Fig. 14 shows the targetless range-averaged beam-Doppler CLEAN mode power.

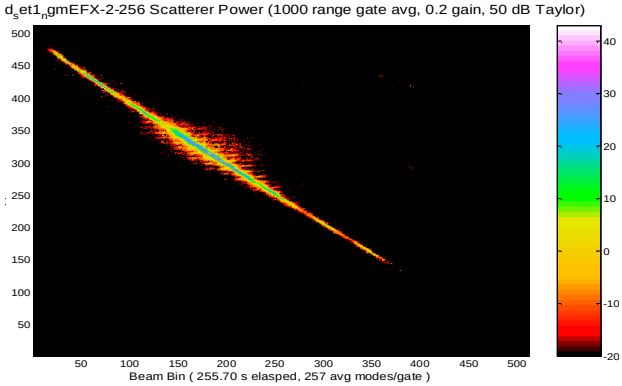

Fig. 14. Targetless range-averaged CLEAN modes.

Note that the SCHISM analysis processing time for the focused datacube decreased from 778 ms to 259 ms per range gate. A more efficient parallel peak finding algorithm was primarily responsible for the improved efficiency. Elimination of the error lobes reduced the previous 321 average modes per range gate to 260.

IMPROVED SCHISM PARAMETER ESTIMATION

The airborne GMTI radar version of the CLEAN algorithm nicely concentrates the modes in beam-Doppler space, but requires an order of magnitude too many modes to prevent lobe leakage. The algorithm can be modified to reduce the number of modes. First, find the frequency of largest peak of the beam-Doppler data FFT and add it to the mode frequency list. Using the listed frequencies and their coherent point spread functions, solve for the mode amplitudes which most closely fit the beam-Doppler FFT of data. Perturb all frequencies and again solve for amplitudes. Retain the perturbations that reduce the modeling error. Repeat for the next error residue peak until power drops below a desired threshold (e.g. noise floor + 13 dB).

Fig. 15 shows the resulting mode parameter fit to the range-averaged beam-Doppler power for the focused datacube, using the usual 13 dB above the noise floor stopping point. The modes are highly concentrated, and more impulsive since only an average of 17 modes is used per range gate. This brute force approach takes an average of 1127 ms to process a single range gate.

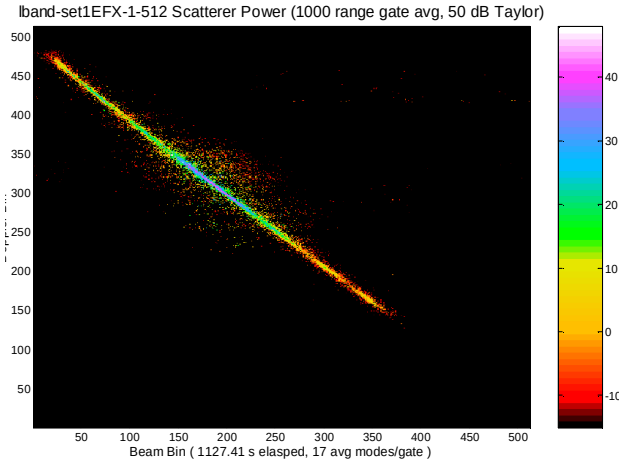

Fig. 15. Mode parameter fit to the range-averaged beam-Doppler power for the focused datacube.

Fig. 16 magnifies the center portion of Fig. 15. Fig. 17 shows the target-only version for comparison to Fig. 16. The model concisely and precisely fits the data, but does not accurately locate all scatterers. Because the clutter is so much stronger than the targets, small model errors can produce many false alarms. Now that fewer modes are used to model each range gate, small errors are much more damaging than before. For the same reason, clutter leakage is less exaggerated by the 1000 range bin average.

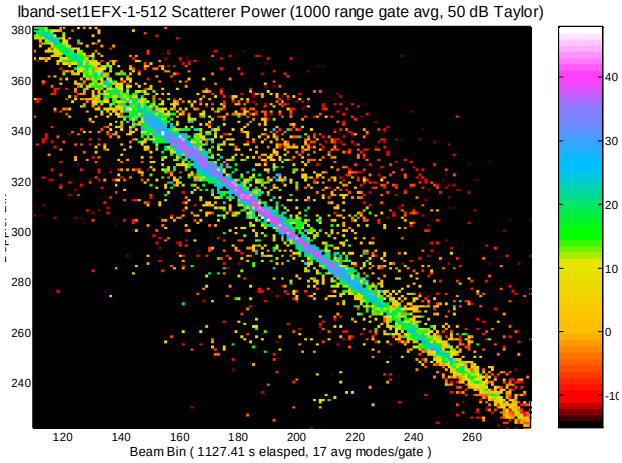

Fig. 16 Detail of mode parameter fit to the range-averaged beam-Doppler power for the focused datacube

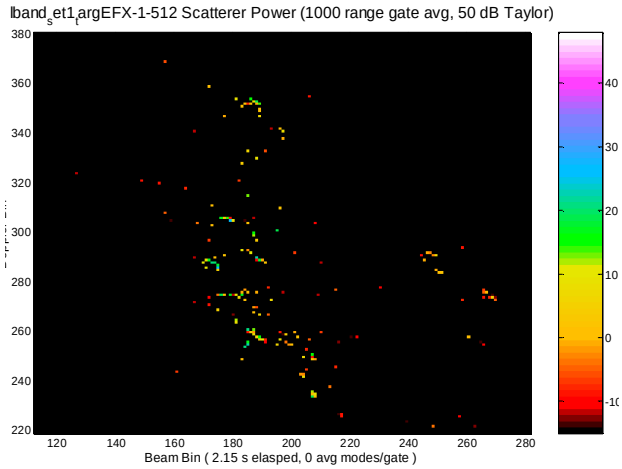

Fig. 17. Target-only parameter fit to the focused datacube beam-Doppler response.

Fig. 18 shows the range-averaged beam-Doppler power of the mode parameter fit to the targetless focused datacube. This plot clearly shows the location of false alarms in beam-Doppler space. Such plots will be invaluable for future algorithm development.

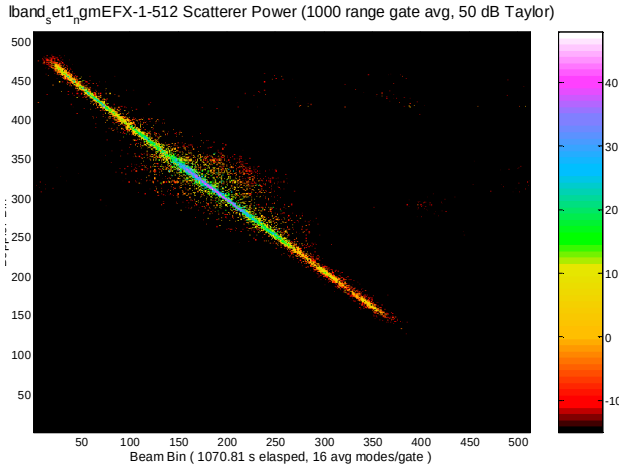

Fig. 18. Targetless mode parameter fit.

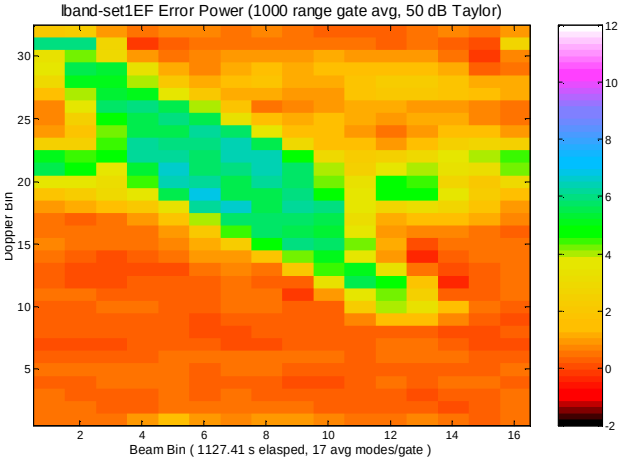

Fig. 19. Mode parameter fit estimation error range-averaged beam-Doppler power for the focused datacube.

Fig. 19 shows the range-averaged beam-Doppler estimation error power for the mode parameter fit to the focused datacube. The peak error was 8 dB even though the algorithm quit at 13 dB above the noise floor. This implies that the mode modeling can be easily extended closer to the noise floor. Also note the valley along the clutter ridge. This indicates a slight broadening of clutter modes.

If the space-time data are not tapered for low FFT sidelobes, peak amplitude estimates are corrupted and the CLEAN algorithm performs poorly. Fig. 20 displays the resulting mode parameter fit to the focused datacube when the data is not tapered. The parameter fit accuracy is actually helped by the added FFT resolution, but the analysis process is slowed to 1606 ms for the average 19 modes per range gate.

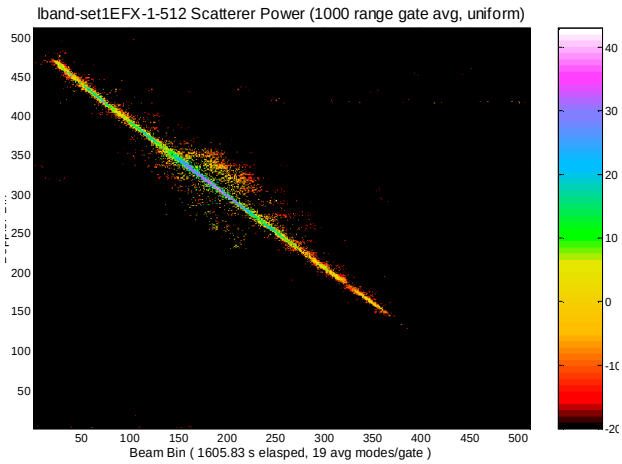

Fig. 20. Mode parameter fit to the focused datacube without the use of a data taper.

Beam-Doppler clutter frequency estimation is tricky because the modes are not orthogonal. One bad mode can corrupt the entire solution, producing many false alarms. The error modes displayed in Fig. 3 can easily be mistaken for mode pairs and vice versa.

QUESTIONS RAISED

The mode parameter fit of (1) to airborne GMTI radar space-time data appears to offer many benefits. Because of the nonlinear nature of beam-Doppler frequency estimation, there are many possible candidate analysis algorithms but not much theory to guide the choice. At this early stage, mode parameter fitting algorithms show promise, but tend to generate many false alarms. On-going efforts strive to build mode fit algorithms that minimize the number of beam-Doppler modes while reducing clutter mode leakage. The availability of target-only and targetless datacubes is invaluable for algorithm development and testing.

Many issues need to be addressed. Is the clutter leakage caused by uncorrected measurement biases or some fundamental limitation of mode analysis? Do analysis procedures yielding fewer modes actually produce "better" results? Are target modes damaged by analysis procedures? Should the analysis be stopped at a higher threshold and used as a preprocessor for standard STAP? How can the analysis algorithm be altered to reduce false alarms? Is recursive high-order cal-on-clutter over successive range bins feasible? Can several CPIs be coherently merged using clutter for alignment? Will SCHISM work on a genuine datacube? Current mode analysis algorithms utilize parallel peak finding to reduce the processing time to under 200 ms per range gate. Is there an even more efficient (hopefully systolic) algorithmic architecture for SCHISM?

REFERENCES

[1] G. R. Legters, "Signal and Clutter as Highly-Independent Structured Modes (SCHISM)", presented at the KASSPER Workshop, MIT Lincoln Laboratory, Lexington, MA, September 23, 2002.

[2] S. Haykin, *Array Signal Processing*. Englewood Cliffs, NJ: Prentice-Hall, 1985, ch. 5, pp. 331-338.

[3] J. R. Guerci, *Space-Time Adaptive Processing for Radar*. Artech House, August 2003.

[4] R. Klemm, *Space-Time Adaptive Processing: Principles and Applications*. London: IEE Press, 1998.

[5] T. K. Moon and W. C. Stirling, *Mathematical Methods and Algorithms for Signal Processing*. Upper Saddle River, NJ: Prentice-Hall, 2000, ch. 13, pp. 391-620.

[6] S. M. Kay, *Fundamentals of Statistical Signal Processing, Vol. 1*. Englewood Cliffs, NJ: Prentice-Hall, 1993, ch. 15, pp. 539-544.

[7] H. Stark and Y. Yang, *Vector Space Projections*, New York, NY: John Wiley & Sons, ch. 3, pp. 91-132.

[8] T. K. Moon and W. C. Stirling, *Mathematical Methods and Algorithms for Signal Processing*. Upper Saddle River, NJ: Prentice-Hall, 2000, ch. 8, pp. 396-399.

[9] T. K. Moon and W. C. Stirling, *Mathematical Methods and Algorithms for Signal Processing*. Upper Saddle River, NJ: Prentice-Hall, 2000, ch. 6, pp. 336-342.

[10] D. E. Dudgeon and R. M. Mersereau, *Multidimensional Digital Signal Processing*. Englewood Cliffs, NJ: Prentice-Hall, 1984, ch. 6, pp. 315-339.

[11] J. S. Bergin and P. M. Techau, "Radar Phenomenology Modeling and High-Fidelity Data Generation", at the KASSPER Workshop, Capitol Hilton, Washington, DC, April 3, 2002.